UNIVERSITY OF OREGON

March 5, 1996

Marilyn J. McKee
AFOSR/PKA
110 Duncan Ave. Room B115
Bolling AFB, DC 20332-8080

RE: Award No. F49620-92-J0384 Final Technical Report

Dear Ms. McKee,

Attached is the Final Technical Report for Award No. F49620-92-J-0384 (Ginsberg)
as you requested.

I can be contacted at (541) 346-3146 for more information if needed.

Sincerely,

Gary Chaffins,
Director

Enc.

c: 235130

19960502 032

OFFICE OF GRANTS AND CONTRACTS 541
5215 University of Oregon · Eugene OR 97403-5215 · Telephone (503) 346-3146

*An equal-opportunity, affirmative-action institution committed to cultural diversity
and compliance with the Americans with Disabilities Act*

REPORT DOCUMENTATION PAGE

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	8 Sep 95	Final Report: 1 Jul 92 - 30 Jun 95

4. TITLE AND SUBTITLE

Real Time Control of Reasoning

5. FUNDING NUMBERS

Project 61102F

Task 2304/GS

6. AUTHOR(S)

Matthew L. Ginsberg

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

University of Oregon
Eugene, OR 97403

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)

AFOSR/NM
110 Duncan Avenue Suite B115
Bolling AFB, DC 20332-0001

10. SPONSORING / MONITORING AGENCY REPORT NUMBER

F49620-92-J-0384

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION / AVAILABILITY STATEMENT

APPROVED FOR PUBLIC RELEASE;
DISTRIBUTION UNLIMITED

12b. DISTRIBUTION CODE

13. ABSTRACT (Maximum 200 words)

This award led to six major technical advances during the contract period. Several of these (approximate planning, dynamic backtracking and limited discrepancy search) promise to substantially change the way various AI subcommunities solve problems.

Approximate planning formalizes an approach to planning that, instead of being correct (every plan returned achieves the goal) and complete (all such plans are returned), is *approximately* correct and complete, in that most plans returned achieve the goal and that most such plans are returned. The cached plans used by case-based planners are best thought of as approximate as opposed to exact, and the approximate approach can be used to attack planning subgoals separately and then combine the plans generated to produce a plan for the original goal. The computational benefits of working with subgoals separately have long been recognized, but attempts to do so using correct and complete planners have failed.

Dynamic backtracking and limited discrepancy search are new approaches to solving constraint-satisfaction problems of the sort that arise in scheduling and other applications. Both allow the flexibility of "lateral" movements in the search space, enabling far more efficient searches and leading to significant performance improvements in systems solving realistic problems.

14. SUBJECT TERMS

15. NUMBER OF PAGES

6

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
UNCL	UNCL	UNCL	

Real-Time Control of Reasoning

Matthew L. Ginsberg

CIRL
1269 University of Oregon
Eugene, OR 97403
ginsberg@cirl.uoregon.edu

Abstract

This award led to six major technical advances during the contract period. Several of these (approximate planning, dynamic backtracking and limited discrepancy search) promise to substantially change the way various AI subcommunities solve problems.

Approximate planning formalizes an approach to planning that, instead of being correct (every plan returned achieves the goal) and complete (all such plans are returned), is *approximately* correct and complete, in that most plans returned achieve the goal and that most such plans are returned. The cached plans used by case-based planners are best thought of as approximate as opposed to exact, and the approximate approach can be used to attack planning subgoals separately and then combine the plans generated to produce a plan for the original goal. The computational benefits of working with subgoals separately have long been recognized, but attempts to do so using correct and complete planners have failed.

Dynamic backtracking and limited discrepancy search are new approaches to solving constraint-satisfaction problems of the sort that arise in scheduling and other applications. Both allow the flexibility of "lateral" movements in the search space, enabling far more efficient searches and leading to significant performance improvements in systems solving realistic problems.

1 Technical progress

1.1 Anytime reasoning

Problem: Declarative problem-solving methods cannot be interrupted; they run for an extended amount of time (since they are typically solving intractable problems) and then return an answer. This is inappropriate in real-time applications.

Previous solution attempts: Many. Typical are Dean and Boddy's proposal of a framework but no real mechanism [10] and Elgot-Drapkin and Perlis's description of a brittle mechanism that depends critically on the choice of language used in any particular application and is likely to be unstable in realistic systems [11].

Key idea: Use modal operators already present in the declarative database as markers for points at which inference can be suspended and a meaningful partial answer returned. Computation can be resumed if better answers are needed and more time is available for analysis. Because modal operators have well-defined meanings independent of language choice, this overall approach is both more robust and more stable than previous work.

Status: Published in the *Journal of Automated Reasoning* [15].

Impact: In addition to its intrinsic merits, this work is enabling technology. It led directly, for example, to the development of approximate planning. It should also lead to a better general methodology for fielding declarative systems in real-time environments.

1.2 Approximate planning

Problem: Existing generative planning tools are unable to debug complex plans in the face of new information or execution difficulties. They cannot reason in real time, do not support such crucial objectives as economy of force, serendipity, component reuse and parallelization, and cannot be used in conjunction with human input (so-called *mixed-initiative* planning).

Previous solution attempts: The planning community recognizes that the above difficulties can largely be addressed by a system capable of solving subgoals separately and then merging the results. Unfortunately, no principled planning system functions in this fashion. Tate's O-PLAN system separates conjuncts but there is no justification for the mechanism used [9]. Kambhampati discusses the need for plan debugging but is unable to address separate conjuncts individually [22]. Plan merging is discussed by Foulser et.al. but no method is presented for incorporating the technique directly into a planner [12]. The best known formal methods make no use of plan merging whatsoever [19, 26].

Key idea: It is possible to plan for conjuncts separately if (and only if) one allows for the possibility that the resulting subplans may interfere when merged. The approximate

1

planning work formalizes this idea and develops precise analyses that bound the importance of the inaccuracies the merge may introduce. It also provides a mechanism whereby the inaccuracies can be reduced if additional computational resources are available.

Status: To appear in the academic literature shortly [14].

Impact: One reviewer of the above paper described it as the most important contribution to generative planning in almost a decade. It should substantially change the AI community's approach to planning generally, and contribute to the use of well-founded generative planners in practical applications.

1.3 Dynamic backtracking

Problem: Recording the reason that certain decisions were made can improve the efficiency of tools that solve constraint-satisfaction problems (such as scheduling). Unfortunately, keeping such justifications has required an amount of memory linear in the run time, and therefore exponential in the problem size. This has made justification-based method useless in practice.

Previous solution attempts: None. It was widely believed that no polynomial-space justification-based algorithm could exist.

Key idea: By retaining only those justifications that constrain the current partial solution, an algorithm can be developed that is guaranteed to find a solution if one exists while using only a polynomial amount of memory.

Status: Published in the academic literature [13, 18].

Impact: Several other authors are now pursuing this line of work. Honeywell, for example, has used the ideas to improve the speed of one of their scheduling tools by an order of magnitude.

1.4 Multiprobe

Problem: Existing search methods are subject to the "early mistake" problem, where they spend too much time in blind alleys, pursuing early bad decisions that eventually render problems insoluble. This difficulty leads to substantial theoretical obstacles and poses significant practical roadblocks to systems attempting to solve optimization problems in real time.

Previous solution attempts: A variety of techniques have been proposed, including iterative approaches [17] and stochastic methods [25, 31]. The iterative methods typically cannot be applied to Boolean problems (of which scheduling is an instance [32]) and the stochastic methods do not perform well on realistic examples.

Key idea: Multiprobe combines the most successful nonsystematic approach with a limited amount of systematic search at each node. The stochastic nature of the search overcomes the early mistakes problem; the limited search at each node overcomes the fact that existing nonsystematic methods often overlook nearby solutions to the problems they are trying to solve.

Status: Experimental results on job shop scheduling have shown that a domain-independent implementation of multiprobe performs comparably to previous methods that have been carefully tuned to exploit specific properties of toy domains (e.g., linear objective functions and the absence of state space resource constraints). The work has been reported in a Stanford Ph.D. thesis [20].

Impact: Because multiprobe does not need to be modified to exploit specific properties of the domain in question, it is far more robust than earlier approaches. It should therefore generate immediate improvements in the ability of fielded systems to address problems that do not satisfy the artificial conditions that other researchers have used to tune their algorithms.

1.5 Limited discrepancy search

Problem: Like multiprobe, limited discrepancy search is a technique for overcoming the early mistake problem. It attempts to combine heuristic information about the choices to be made at each node with a search methodology that moves rapidly from one portion of the search space to another.

Previous solution attempts: Existing systematic search methods are completely hamstrung by their need to search the space in some predetermined order. It is clear, however, that a more effective approach would be to search the space in a "heuristic" order, examining the nodes in order of the prior probability that they solve the problem in question.

Key idea: Limited discrepancy does indeed examine the entire search space in heuristic order. The essential idea is to rank the fringe nodes by their heuristic merit alone and to then use any of a variety of existing best-first search techniques to search the fringe of the tree appropriately. Korf's IDA* technique [23] is the one currently used.

Status: Published in Harvey's thesis [20] and IJCAI-95 [21]. In addition, limited discrepancy search has been applied to a realistic problem relevant to aircraft manufacture (provided to CIRL by an employee of McDonnell Douglas Aerospace), and led to substantial reductions in the length of the shortest known solution schedule.

Impact: As with multiprobe, limited discrepancy search does not need to be modified to exploit specific properties of the domain in question. Other researchers have already applied the technique to the toy domain of number partitioning [24], and work is under

way to apply it to TPFDD scheduling and the constraint satisfaction problems arising in the CTEMS air operations scheduling tool.

1.6 Partition search

Problem: Existing adversarial search methods need to analyze every possible opposing action separately, even though most of these actions can typically be defeated in the same fashion.

Previous solution attempts: Adelson-Velskiy et.al.'s *method of analogies* [1]. This approach is not used in practice because its formal underpinnings are weak and it fails to combine opposing actions that are similar but not identical.

Key idea: Use dependency techniques to automatically identify classes of positions with identical outcomes, and then back these classes up through the search tree. By doing so, the transposition tables maintained by adversary search programs can be made far more effective, leading to an increased number of cache table hits and substantial reductions in overall run time.

Status: "Partition search" is available as a CIRL technical report [16]. In the toy domain of bridge, the method leads to order-of-magnitude savings over the best competing approach.

Impact: Unknown. The consensus of the members of the game-playing community to whom the method has been described is that it will lead to large computational savings. There are straightforward arguments that applications to adversarial military situations should also be possible, but this is difficult to test because we have been unable to obtain realistic adversarial planning data.

References

† = publication acknowledges support from this contract

[1] G. Adelson-Velskiy, V. Arlazarov, and M. Donskoy. Some methods of controlling the tree search in chess programs. *Artificial Intelligence*, 6:361–371, 1975.

† [2] A. B. Baker. The hazards of fancy backtracking. In *Proceedings of the Twelfth National Conference on Artificial Intelligence*, 1994.

† [3] T. Bedrax-Weiss and L. Bertossi. An assessment of Reiter's general monotonic solution to the frame problem and its applicability to database updates. In *Proceedings of the ECAI-94 Workshop on Logic and Change*, pages 5–18, Amsterdam, The Netherlands, Aug. 1994.

† [4] T. Bedrax-Weiss and L. Bertossi. An underlying semantics for Reiter's solution to the frame problem. In *Proceedings of the XII Brazilian Symposium on AI (SBIA-95)*. Springer-Verlag, Oct. 1995. To Appear.

† [5] J. M. Crawford and L. D. Auton. Experimental results on the crossover point in random 3SAT. *Artificial Intelligence*, 1995. To appear.

† [6] J. M. Crawford and A. B. Baker. Experimental results on the application of satisfiability algorithms to scheduling problems. In *Proceedings of the Twelfth National Conference on Artificial Intelligence*, 1994.

† [7] J. M. Crawford, D. Dvorak, D. Litman, A. Mishra, and P. Patel-Schneider. Path-based production rules. In *Workshop on Embedded Production Systems in Object-Oriented Languages at The Ninth Annual ACM Conference on Object-Oriented Programming Systems, Languages and Application*, October 1994.

† [8] J. M. Crawford and D. Etherington. Observations on observations in action theories. In *Proceedings of the 1995 AAAI Spring Symposium on Extending Theories of Action*. American Association for Artificial Intelligence, March 1995.

[9] K. Currie and A. Tate. O-plan: The open planning architecture. *Artificial Intelligence*, 52:49–86, 1991.

[10] T. Dean and M. Boddy. An analysis of time-dependent planning. In *Proceedings of the Seventh National Conference on Artificial Intelligence*, pages 49–54, 1988.

[11] J. J. Elgot-Drapkin and D. Perlis. Reasoning in time I: Basic concepts. *J. Expt. Theor. Artif. Intell.*, 2:75–98, 1990.

[12] D. E. Foulser, M. Li, and Q. Yang. Theory and algorithms for plan merging. *Artificial Intelligence*, 57:143–181, 1992.

† [13] M. L. Ginsberg. Dynamic backtracking. *Journal of Artificial Intelligence Research*, 1:25–46, 1993.

† [14] M. L. Ginsberg. Approximate planning. *Artificial Intelligence*, 1995.

† [15] M. L. Ginsberg. Modality and interrupts. *J. Automated Reasoning*, 14:43–91, 1995.

† [16] M. L. Ginsberg. Partition search. Technical report, CIRL, University of Oregon, Eugene, OR, 1995.

† [17] M. L. Ginsberg and W. D. Harvey. Iterative broadening. *Artificial Intelligence*, 55:367–383, 1992.

† [18] M. L. Ginsberg and D. A. McAllester. GSAT and dynamic backtracking. In *Proceedings of the Fourth International Conference on Principles of Knowledge Representation and Reasoning*, Bonn, Germany, 1994.

[19] S. Hanks and D. S. Weld. Systematic adaptation for case-based planning. In *Artificial Intelligence Planning Systems: Proceedings of the First International Conference*, pages 96–105. Morgan Kaufmann, 1992.

† [20] W. D. Harvey. *Nonsystematic Backtracking Search*. PhD thesis, Stanford University, Stanford, CA, 1995.

† [21] W. D. Harvey and M. L. Ginsberg. Limited discrepancy search. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*, pages 607–613, 1995.

[22] S. Kambhampati. Design tradeoffs in partial order (plan space) planning. In *Artificial Intelligence Planning Systems: Proceedings of the Second International Conference*, pages 92–87. AAAI Press, 1994.

[23] R. E. Korf. Depth-first iterative deepening: An optimal admissible tree search. *Artificial Intelligence*, 27:97–109, 1985.

[24] R. E. Korf. Improved limited discrepancy search. Technical report, UCLA, Los Angeles, CA, 1995.

[25] P. Langley. Systematic and nonsystematic search strategies. In *Artificial Intelligence Planning Systems: Proceedings of the First International Conference*, pages 145–152. Morgan Kaufmann, 1992.

[26] D. McAllester and D. Rosenblitt. Systematic nonlinear planning. In *Proceedings of the Ninth National Conference on Artificial Intelligence*, pages 634–639, 1991.

† [27] J. C. Pemberton. *Incremental Search Methods for Real-Time Decision Making*. PhD thesis, UCLA, 1995.

† [28] J. C. Pemberton. k-best: A new algorithm for making real-time decisions. In *Working Notes of the IJCAI-95 Workshop on Anytime Algorithms and Deliberation Scheduling*, pages 65–72, 1995.

† [29] J. C. Pemberton. k-best: A new method for real-time decision making. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence (IJCAI-95)*, pages 227–233. Morgan Kaufmann, August 1995.

† [30] R. Schrag and J. M. Crawford. Implicates and prime implicates in random 3SAT. *Artificial Intelligence*, 1995. To appear.

[31] B. Selman, H. Levesque, and D. Mitchell. A new method for solving hard satisfiability problems. In *Proceedings of the Tenth National Conference on Artificial Intelligence*, pages 440–446, 1992.

[32] S. F. Smith and C.-C. Cheng. Slack-based heuristics for constraint satisfaction scheduling. In *Proceedings of the Eleventh National Conference on Artificial Intelligence*, pages 139–144, 1993.

PROPOSAL NUMBER: 92NM295 (2295f/p15)
TYPE SUBMISSION: Final
INST. CONTROL NUMBER: F49620-92-J-0384
INSTITUTION: Univ of Oregon
P.I. NAME: Professor Matthew Ginsberg
INVENTION IND: NONE
PROJECT/TASK: 2304 GS
PROGRAM MANAGER: DR ABRAHAM WAKSMAN

****

1. Objective:
@36.1@UNARR. ------   OBJT.
!! The aim of this research is to develop new ideas on the effective control
of reasoning in real-time and other declarative AI systems. It is propose to
do this by building on the anytime work that has already been supported by
AFOSR as part of the IRTPS effort. Specifically: 1. Developing a formal
framework in which modal operators are viewed not only as interruption
markers but as control markers as well. 2. Showing that existing
control heuristics can be easily recast in this framework. 3. Developing
and exploring new and more general heuristics that can best be expressed
only in this new framework.

2. Approach:
@37.1@UNARR. ------   APPR.
$$ In the earlier IRTPS effort under an AFOSR grant, it was suggested that
modal operators, in addition to their well-known semantic role in declarative
systems, also played the role of interruption markers, giving semantic
indications of those points at which the analysis of a reasoning system could
be suspended and a partial answer returned. In this proposed effort, this
idea will be pursued further, showing that modal operators also serve to
mark points at which control information can be brought to bear. These two
ideas are closely coupled; it is reasonable to expect that there should be
some identity between the points at which a system can be interrupted to
return a partial answer and the points at which it can be interrupted to
apply control information.

3 Progress:
@38.1@UNARR. ------ PROG.
See report for abstract